VERITAS

NetBackup Flex Scale

Data protection that's highly secure and resilient, simple to scale, and easy to operate.

White Paper | March 2022

Contents

Version	Details
Jan 2021	Initial release.
Mar 2021	Added references. Updated for NetBackup Flex Scale 1.3.1.
Nov 2021	Updated for NetBackup Flex Scale 2.1.
Mar 2022	Updated for NetBackup Flex Scale 3.0.

Executive Summary

In today's data-driven world, ensuring business data is well protected is of paramount importance. The continuous growth of data has created a constant and increasingly pervasive risk of malware infecting that data. In this reality, data protection solutions must be capable of effectively defending against malicious data alteration or deletion without sacrificing scale or automation requirements.

In the next-generation data center, IT departments require a data protection solution that can:

- Protect the full range of workloads across an enterprise
- Provide highly scalable space-optimized capacity that can protect more data with shorter backup and faster recovery windows
- Deliver backup/recovery performance that scales as protection needs increase over time
- Ensure recovery of business-critical data at any scale
- Provide an effective defense and mitigation for ransomware attacks with data and infrastructure immutability

Veritas NetBackup™ Flex Scale is a highly flexible, hyperconverged data protection solution that applies the proven power of the NetBackup architecture to the needs of the next-generation data center. It serves an organization's urgent requirements for enterprise resiliency, simplified operations, and a highly utilized, pay-as-you-grow, scale-out architecture.

A NetBackup Flex Scale deployment begins with a smaller initial configuration that meets existing performance and capacity requirements. As these requirements increase, automated processes grow the cluster to add more capacity and backup/recovery performance. Depending on the model, its total usable capacity can be scaled up to 1.8 PB (1.6 TiB) of usable space within a single domain on local storage, with the option to add an extended cloud-based storage tier or export stored data to be written to tape for long-term retention (LTR) purposes. Because it uses NetBackup's advanced deduplication engine, data savings are high locally, in transit during replication, and in long-term, cloud-tier storage.

NetBackup Flex Scale was designed with a focus on simplifying operations. It provides streamlined management with orchestrated automation of operations, reducing operational risk and overall management costs. The combination of NetBackup Flex Scale infrastructure management with the familiar NetBackup user interface (UI) eliminates the learning curve for administrators.

Technical Overview

When designing a data protection solution, organizations have a variety of requirements that require flexible deployment options for their different data centers.

Veritas understands that data protection can't be managed with a one-size-fits-all solution. That's why NetBackup offers a wide variety of deployment options, providing the flexibility to implement a data protection solution that best suits the specific needs of data in remote offices, in the cloud, or in a central data center that requires the ability to scale both capacity and compute resources.

In NetBackup's case, these deployment options are not simply variations on the same essential architecture but extend it to accommodate a huge variety of data protection requirements.

The newest deployment option, NetBackup Flex Scale, accommodates those requirements for operational simplicity and scalability within a single domain with ransomware resiliency and enterprise-grade availability while leveraging the stability and time-tested power of NetBackup itself.

Operational Simplicity

Reducing management costs and complexity has been a long-standing design focus at Veritas. The management of NetBackup Flex Scale is streamlined, including automated NetBackup and infrastructure operations. To make things easy to manage, it starts with a single UI that combines the same familiar NetBackup interface with NetBackup Flex Scale infrastructure details. (See Figure 1.)

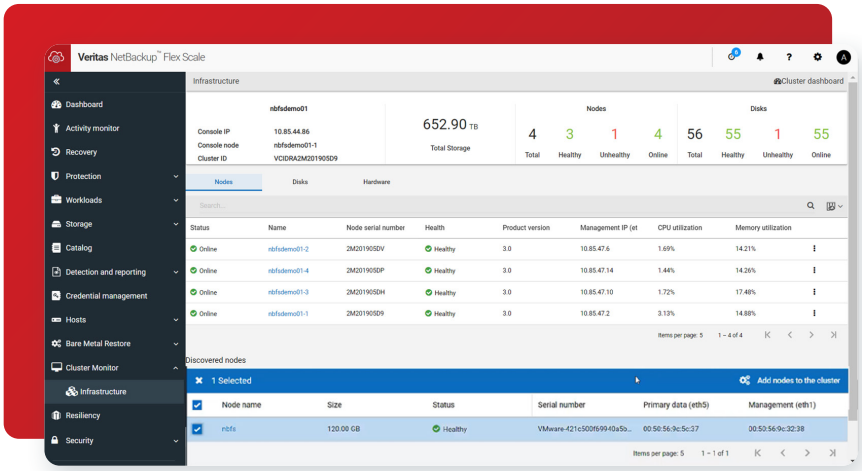

Figure 1. NetBackup Flex Scale displays infrastructure information in the familiar NetBackup UI.

Operational simplicity is also prevalent throughout the management lifecycle, starting with the automated process involved in the initial deployment, initial backup administration, and onwards through common infrastructure management tasks such as cluster scaling and recovery operations.

There are many automated processes built directly into NetBackup Flex Scale (see Figure 2).

Some examples of those processes include:

- An initial configuration with an automated process that installs and configures the nodes, network, cluster, and NetBackup software in a highly resilient configuration without additional user input
- Cluster scaling that auto-detects new nodes and background tasks that shift data across the cluster to increase performance, balance capacity, and allow more NetBackup service instances to accept data protection jobs

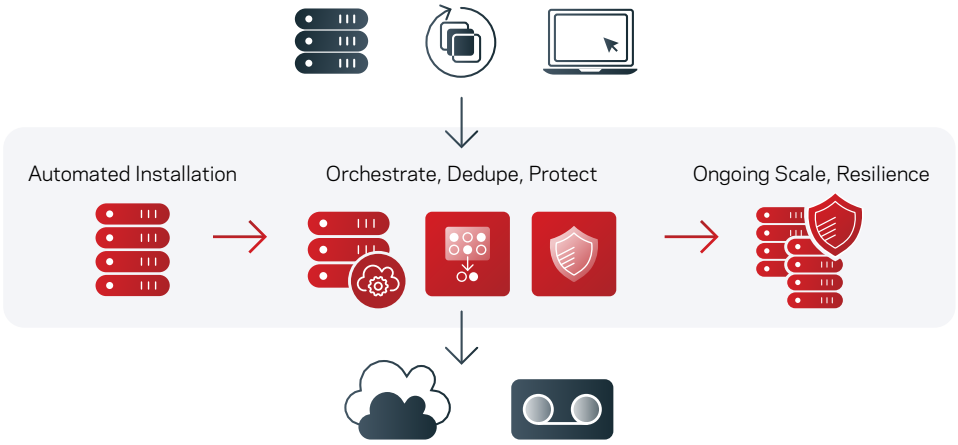

Figure 2. Automated processes in NetBackup Flex Scale that simplify operational tasks.

- Integrated intelligent load balancing that distributes backup and recovery jobs to provide the optimal distribution across the cluster
- Technology refreshes using an automated process that replaces the data and configuration from an existing node with a new/replacement node without the need to manually perform catalog or primary server migration operations
- A lockdown mode process that can be used to increase the cluster's operational and data security in addition to immutable or write once, read many (WORM) storage
- An automated upgrade and patching process that allows software updates to be acquired directly from the Veritas Service Operations and Readiness Tools (SORT) site or uploaded directly to the NetBackup Flex Scale interface, before NetBackup Flex Scale handles the process of distributing the data to cluster nodes and performing a rolling upgrade so scheduled data protection jobs can continue

Automated Installation and Provisioning

NetBackup Flex Scale's built-in automation makes deployment a straightforward and consistent process, with input needed only during a single initial step to become fully operational.

The first step is to provision data center space and network access for the NetBackup Flex Scale Appliances, physically install them in a rack, and connect the appliances' power and network. Each appliance has a pair of 10/25-GB Ethernet ports for private cluster traffic and another pair of 10/25-GB Ethernet ports for client-facing data protection traffic.

All that's necessary to begin the deployment process is to make the management port of one of the nodes accessible, typically using the system console. You access the NetBackup Flex Scale installer on this interface, and it handles the rest of the orchestration for the install process. Initially, it discovers the other unconfigured nodes by searching the shared network context for the presence of their private network interfaces and performs a health check to ensure they are ready for the upcoming procedure.

The only step you need to take is to inform the NetBackup Flex Scale deployment of its configuration, either using API calls or the integrated web UI. You can provide this initial configuration by manually entering the configuration details or by providing a human-readable unattended installation file with predefined configuration details such as networking and naming specifics.

No further attention is needed once the deployment process begins. NetBackup Flex Scale distributes the configuration to the appliances, configures the nodes, creates a NetBackup cluster with a distributed, deduplicated storage pool, configures NetBackup, and starts the containerized services.

After deployment is complete, NetBackup is fully configured and ready to discover your assets and create protection policies. There are also API calls and UI capabilities available to administer, monitor, and update the physical appliance hardware.

Deployment Options

During deployment, NetBackup Flex Scale gives the flexibility to be deployed as either a turnkey, drop-in data protection solution and as an instantly ready, scale-out capable expansion of media service and storage capability for an already-deployed NetBackup infrastructure. This flexibility lets you introduce NetBackup Flex Scale without changing your existing backup architecture. (See Figure 3.)

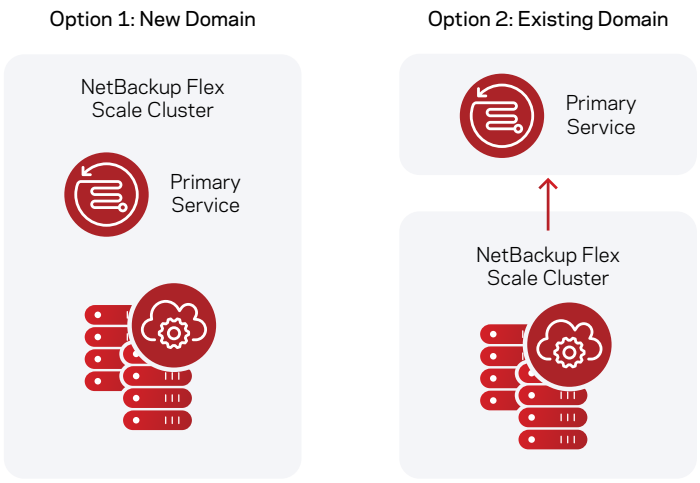

Figure 3. NetBackup Flex Scale deployment options.

You can choose a media-server-only deployment which allows you to specify an external Primary server (service) instead of deploying NetBackup Flex Scale as a new domain with its own Primary server. This deployment option has the effect of integrating the scale-out services provided by NetBackup Flex Scale with an existing NetBackup domain infrastructure

It should be noted that this deployment option prevents you from leveraging the scaling, resiliency, and availability features of the NetBackup Flex Scale Primary service and catalog, including automated expansion of the catalog and site failover; this instead will be managed off-cluster by the existing infrastructure

Automated Scheduling of NetBackup Jobs

Immediately after deployment, the NetBackup Flex Scale deployment presents a new, fully functional NetBackup domain, with services available, storage configured, administrative users set up, and each node configured to contribute its storage and compute resources.

As a result, an administrative user specified during setup is immediately capable of discovering assets and protecting workloads using protection plans that automate the scheduling of protection jobs. These jobs will automatically be intelligently distributed across the cluster, maximizing throughput by using all available resources within the cluster.

Increasing Cluster Size

A data protection solution will inevitably need to grow to accommodate its growing data footprint. This growth includes scaling both system resources and storage pool capacity. NetBackup Flex Scale's high level of automation makes adding new resources to the cluster a simple and repeatable procedure. You follow the same steps to add a single node or multiple nodes at once, expanding from a minimum of 4 nodes to a maximum of 16 nodes.

The process of adding a node to an existing NetBackup Flex Scale cluster is as simple as powering the appliance on in the same network context as the existing deployment, initiating its discovery using an API call or the NetBackup Flex Scale UI, and providing the network addresses and names it needs to begin expansion.

NetBackup Flex Scale automates the scale-out operation, which involves:

- Configuring the node and its network
- Joining the cluster
- Rebalancing data in an optimized fashion across the cluster
- Starting additional NetBackup services

Once a cluster expansion has begun, no further action is needed. NetBackup jobs continue uninterrupted. NetBackup Flex Scale will efficiently transfer the minimal amount of data required across the private cluster network to the new node to achieve a balanced and resilient data pool. Afterwards, the new node(s) will be included in the intelligent balancing of NetBackup jobs, and will also begin intelligently routing data protection jobs automatically to the newly added cluster resources.

There are no changes needed on the NetBackup client side before or afterwards: the same network endpoints are used regardless of how large the NetBackup Flex Scale cluster grows.

Replacing Nodes

At some point, an administrator may have to replace a node in a cluster, either due to hardware failure or to use a newer model of the NetBackup Flex Scale Appliance. This procedure is also orchestrated by a NetBackup Flex Scale internal process, so services running on other nodes remain uninterrupted for the duration.

Even if the node being replaced is running the Primary service, the automated process will automatically restart the service on another node, and because the catalog is mirrored across all the nodes in the cluster, no migration of catalog data is required before or after the replacement.

Enterprise Resiliency

NetBackup Flex Scale takes a resilient and distributed approach to ongoing operations by eliminating multiple points of failure in hardware and using its clustering components to remain continuously available.

Appliance and Storage Availability

NetBackup Flex Scale's highly available design with clustered services and distributed data allows for a fully orchestrated service restart or redeployment. The erasure-coded disk pool can continue ingesting and restoring data even while sustaining the loss of multiple disks, an entire node, or both at once. As you add more appliances to the NetBackup Flex Scale deployment, the potential ratio of sustained disk and node losses to intact ones increases and the cluster automatically becomes more resilient.

Cluster Size		Disk Pool Failure Tolerance	
	Disks	Nodes + Disks	Nodes
4–5	Any 4 HDDs	1 Node + any HDD	1
6–12	Any 4 HDDs	1 Node + any 2 HDDs	2
12–16	Any 4 HDDs	1 Node + any 3 HDDs	2

For NetBackup Flex Scale's resiliency purposes, the loss of SSDs backing the NetBackup catalog volume or HDDs backing the deduplicated data pool have a different impact because the catalog volume employs snapshots, mirroring, and ongoing catalog self-backups to keep multiple copies available. HDDs participate in a single, cluster-wide disk pool, but the larger number of devices means it can sustain many failures and continue to operate as the cluster widens.

When deployed with four or five nodes, NetBackup Flex Scale can continue to operate uninterrupted after losing one appliance, plus another disk anywhere in the deduplicated data pool, or any four individual disks anywhere in the data pool. After deploying six appliances in a cluster, NetBackup Flex Scale's redundancy level increases and it can suffer the loss of two separate appliances, or any appliance and an additional two (with 6–10 nodes) or three (with 11–16 nodes) disks from the deduplicated disk pool.

In operation, nodes have an additional resiliency characteristic: The separation of data between nodes for the purpose of resiliency guarantees a given node can only have a given ratio of the total proportion needed to reconstruct it. For example, with 16 nodes, if 4 nodes contain a single failed disk, they can continue to fail disks to a maximum of 48 without losing data.

As a final potential method of increasing data resiliency, NetBackup Flex Scale can also send ingested data to be physically copied onto tape, allowing storage lifecycle policies (SLPs) to include long-term retention (LTR) or off-site data storage for compliance or regulatory reasons.

Site-level Availability

You can also deploy NetBackup Flex Scale in a site-redundant fashion, with two clusters at different sites, to achieve site-level redundancy and availability even in a disaster recovery (DR) scenario. There are two common approaches available to NetBackup Flex Scale:

- Single-domain, dual-site—Both sites contain a NetBackup Flex Scale cluster segment. Each site's NetBackup Flex Scale nodes host media and storage services. Clients may be directed to the media services present at either site's nodes as needed. Deduplicated backup data is replicated between sites using SLPs and optimized duplication, and NetBackup catalog data and Primary service data is replicated between sites using volume replication. In the event of a site failover, the cluster's Primary service restarts at the remote site to immediately make mass recovery and continued protection available.

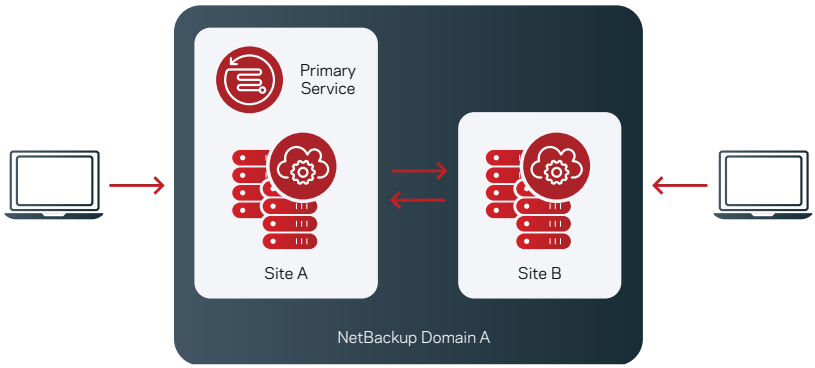

- Dual-domain, dual-site—Each site hosts its own NetBackup domain as a Primary, with both Site A and Site B effectively transmitting data to its partner site. In this case, Site B could be another NetBackup Flex Scale cluster or a different type of NetBackup deployment. The two sites actively receive data to protect from clients and transmit it to the partner site, so the replicated data is available in the other domain via Auto Image Replication (AIR) as with a single-domain deployment. Each side of the cluster can consist of the maximum number of cluster nodes (up to 16 nodes).

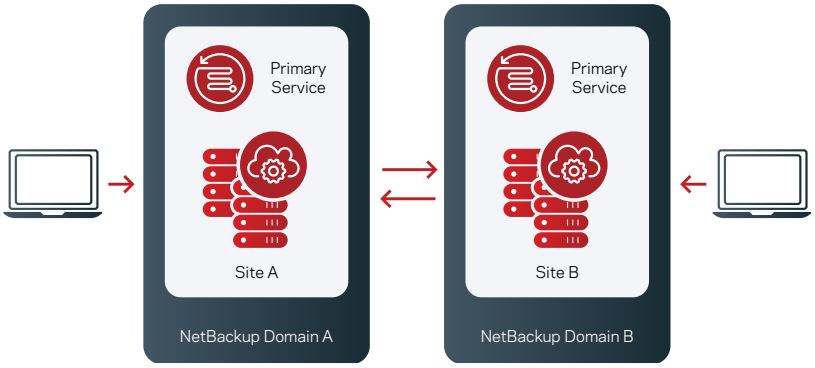

Data Resiliency

A NetBackup Flex Scale deployment has two relevant storage areas (and two potential adjunct areas):

- Hard drives are used to protect deduplicated client data in a highly resilient deduplicated pool. All the hard drives are configured in a clustered file system that is presented as a single storage pool to NetBackup.

  This data is stored in an erasure-coded format with a data-to-parity ratio of 8:4, yielding a storage efficiency ratio of 67 percent

  All data stored in this pool is deduplicated inline, resulting in a usable capacity that is significantly larger than the physical capacity available. Overall, usable capacity varies depending on the specific workloads being protected and their rate of change.

- NVMe flash storage is used for storing and protecting the NetBackup catalog and backup metadata. This storage is configured in a clustered file system that is provisioned with triple mirrors. The capacity will scale automatically as the cluster size increases.
- As adjunct storage, NetBackup Flex Scale can transparently extend its storage pool and write deduplicated backup data directly onto S3-compatible object storage that acts as an extension of the deduplicated data pool.
- Finally, NetBackup Flex Scale can use a media server—a NetBackup Appliance or Flex Appliance or installed on compatible hardware—joined to the same domain with attached tape resources to execute an SLP that writes stored data out to tape for LTR or air-gapped resilience purposes.

Data and Infrastructure Immutability

You can configure NetBackup Flex Scale to provide both its deduplicated data store and its operational infrastructure with immutability properties to guard against ransomware attacks, ensure compliance with legal or organizational data retention requirements, or both. When in this lockdown mode of operation, NetBackup Flex Scale provides an in-depth approach to protection against the malicious deletion or alteration of data—first by protecting inbound data, then by enforcing its immutability, and finally by providing the resilience and mass recoverability enabled by a NetBackup scale-out architecture.

Leveraging a Containerized Architecture for In-depth Security

A service architecture based on containers is inherently more secure due to the isolated nature of container resource allocation and namespaces and their logically separated configurations. Veritas has taken advantage of these attributes to implement a container-based NetBackup architecture running on NetBackup Flex Scale that takes an in-depth, multi-level approach to preventing any unauthorized access to system data or unauthorized use of system resources. (See Figure 4.)

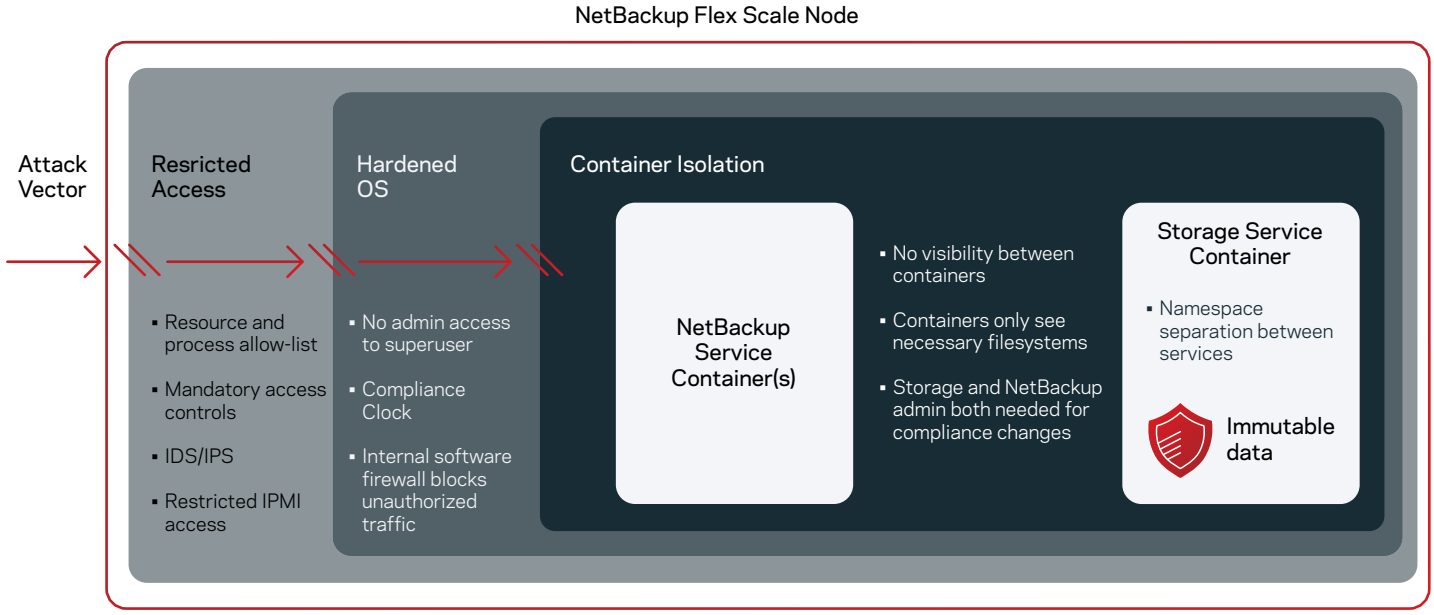

Figure 4. The multi-level strategy NetBackup Flex Scale uses to prevent unauthorized access or use.

- The service containers are designed to eliminate the need for elevated system privileges; the containers themselves are given permissions (using SELinux mandatory access controls [MAC]) only to the system resources required for their continuing operation, preventing them from becoming a method of unauthorized privilege elevation. The container environment will refuse to execute or allocate resources to programs not matching a predefined allowed list.
- Container resources are inherently isolated from each other because they operate in separate Linux Control Groups (Cgroups). These Cgroups hierarchically organize access to node storage and network subsystems as well as system resources and prevent processes running in the context of one Cgroup from visibility into the operating context of another.
- Starting a container-based service is subject to security constraints and checks: rather than being read and executed from a node's operating system files, a container's contents are bundled together in binary form and a checksum comparison is carried out before execution to ensure program contents are immutable.

Infrastructure Immutability and Lockdown Mode

A volume configured to contain immutable data for a given retention period is only capable of delivering on the promise of the retention period if sufficient measures are also taken to ensure the systems and software used to store the data cannot be tampered with either by malicious software or by a rogue administrator or insider.

NetBackup Flex Scale adopts a Zero Trust architecture model when in lockdown mode, reducing the ability of even administrators to change the system's operational state or parameters while tightly restricting access to system components (see Figure 5). For example, in lockdown mode, the following additional restrictions apply:

- Data stored with a retention lock is immutable
- Built-in remote access interfaces (intelligent platform management interface [IPMI]) are restricted to non-disruptive operations
  - Single-user mode/rescue mode boot options are restricted
  - Editing OS bootloader parameters is restricted
- Administrator access to the operating system root account is disabled

Custom-defined SELinux MAC are enabled to protect operating system data and compartmentalize access capabilities of NetBackup and storage services. The MAC implementation means that system and service settings can only be changed by trusted applications or scripts, even with system-level privileges. The platform operating system underlying NetBackup Flex Scale and its containerized services has been further customized to remove unnecessary components and services that could be subject to attack.

Lockdown mode also implements a software Compliance Clock that runs independently of configured system time or Network Time Protocol and continuously whenever the NetBackup Flex Scale cluster is operational. In fact, even a cluster shutdown will not advance the Compliance Clock expiration time. Reference to this clock determines if the specified amount of retention time for a backup image has passed, not the system clock or elapsed real time.

Properties	Standard Protection	Immutable Protection	
	◉ Normal	○ Enterprise	○ Compliance
Immutable data support with retention locking	X	✓	✓
Deletion of immutable data before expiry by the Backup Administrator	✓	X	X
Backup image retention lock deletion	—	✓	X
Access to Remote Management Platform	✓	X	X
Appliance Administrator access to node operating system	✓	X	X
Appliance immutability mode upgrade	✓	✓	—
Appliance immutability downgrade	—	X	X
Retention lock extension	—	✓	✓

Figure5. Lockdown mode selection in the NetBackup Flex Scale user interface.

Data Immutability and WORM

A volume of data is described as WORM if it has the capability to store data immutably. When data at rest is immutable, it cannot be deleted or changed and the system will deny any attempts to shorten the retention period, overwrite the data, or remove it by users and system administrators.

When entering lockdown mode, you can choose either Enterprise or Compliance versions of this capability. Enterprise mode permits data written to a WORM-capable volume to be expired in a two-step process (creating an audit log entry to do so). Both an administrator with the Backup Admin role and an administrator with the Appliance Admin role must participate in this process. In Compliance mode, the retention period can never be shortened, nor can lockdown mode be made less strict or exited while any data with a retention period exists in the datastore.

NetBackup Flex Scale runs a distributed storage server that consists of containerized services on each of the participating cluster nodes. As part of configuring lockdown mode, you set the minimum and maximum possible backup retention time for the storage pool.

Subsequently, when creating backup policies to store data in the NetBackup Flex Scale deduplicated data pool you must specify either a retention period within those limits or that no retention period is required. Multiple storage units can write to the same storage server, and they can be separately configured for WORM retention or without it.

Once the storage server has been moved to lockdown mode, and you have specified a minimum and maximum retention period for the storage service, the storage pool is designated as WORM-capable and you can create storage units that have the WORM capability set.

Third-party Validation of Immutability

Veritas has submitted NetBackup Flex Scale for a regulatory compliance assessment evaluation by Cohasset Associates, with the result of certification for the following requirements:

- Securities & Exchange Commission (SEC) in 17 CFR § 240.17a-4(f)
- Commodity Futures Trading Commission (CFTC) in regulation 17 CFR § 1.31(c)-(d)

NetBackup Flex Scale Architecture and Data Protection

The NetBackup Flex Scale architecture serves as a superset as well as a specific instance of scale-out NetBackup architecture (see Figure 6). It automates common data protection operations, but leverages its architecture for other operational tasks that can be complex or time-consuming, including:

- Deployment—Automated configuration of highly secure OS, clustered filesystem and NetBackup services, storage and catalog protection.
- Updates: Seamless updates that automate infrastructure and NetBackup updates

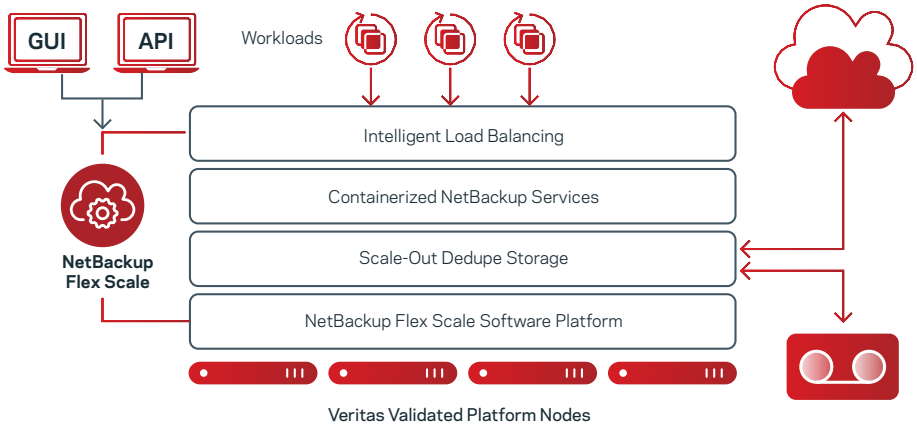

Figure 6. How the NetBackup Flex Scale architecture automates data protection and operational tasks.

- Balancing—Ensures data is logically stored across the cluster to maximize resiliency and performance
- Availability—Services are intelligently monitored to ensure they're always online

This containerized architecture serves to improve NetBackup Flex Scale's efficiency, simplicity, and scalability with minimal resource overhead. It permits a purpose-built architecture to automate common data protection operations while it continues to deduplicate incoming data, commit it to resilient storage, and perform the background tasks of replication and catalog protection.

However, because NetBackup Flex Scale orchestrates the creation of these layers during the deployment process, as described in the section "Automated installation and provisioning," it's not necessary to directly interact with the underlying layers to use them. Instead, they are leveraged using NetBackup Flex Scale's automated resource management or with internal orchestration triggered by an API call in some cases.

Common tasks that are secondary to protecting data from clients but crucial to the ongoing health and availability of the system use NetBackup Flex Scale's purpose-built architecture, which automates these tasks, making them simple for the administrator, and if possible, requiring no input or attention.

Load Balancing

Along with storage-level balancing, NetBackup Flex Scale will also load balance incoming client data protection. It will automatically direct an incoming client's traffic to the last physical node it previously communicated with by default to be fingerprinted, deduplicated, and stored, but relevant factors are capable of overriding this default. For example, if the node is no longer participating in the NetBackup Flex Scale cluster when a node begins a backup during its backup window, or if the node is showing that it is overloaded with other work, the incoming job will be passed to another endpoint. Another possibility is that the cluster has recently scaled out further and previous client affinities are now to be transferred to another physical node to keep front-end throughput distributed.

In short, NetBackup Flex Scale does not simply apply a round-robin approach between available nodes but uses its cluster awareness and intelligence to examine system resource load, number of concurrent active jobs, and last backup location to determine the optimal job distribution.

Performance

When scaling out, NetBackup Flex Scale automates the balancing of data across the erasure-coded disk pool underlying its deduplicated storage. As nodes are added, more devices are added. A larger pool of devices means that not just the pool capacity increases but that more devices have more capacity to service I/O requests, so that concurrency increases with cluster size.

NetBackup Flex Scale also lets an administrator select a performance profile when adding an additional node to an existing cluster, allowing a choice between prioritizing the rebalancing of data onto newly added storage or prioritizing I/O traffic for ongoing backup and restore jobs.

Finally, NetBackup Flex Scale maintains both performance dashboard graphs in its user interface and a highly available API endpoint to service orchestration requests. Queries can be sent for performance statistics and to monitor the status and duration of data protection jobs.

Recovery

NetBackup Flex Scale automates the repair, reprotection, and recovery of its underlying data volumes without user intervention (except to physically replace failed hardware, when necessary). If a node is offline for any amount of time, its offline disks will be populated with the correct data in the background once it returns to service. If a disk actually fails and is replaced by using erasure coding, the cluster's disk pool assists in recalculating parity or data to re-integrate the replacement disk into the pool.

NetBackup Flex Scale also automates the replacement of a node after a failure. If the node will not be capable of returning to the cluster, NetBackup Flex Scale will discover an unconfigured node and assign it the old node's place in the cluster while simultaneously reconstructing missing data onto its disks.

Because NetBackup Flex Scale is always deployed in a clustered configuration with multiple nodes, it also attempts to assist in its own recovery by automating common support tasks such as collecting logs from each NetBackup Flex Scale node in the cluster for troubleshooting setup issues or assisting Veritas Support. The nodes are also independently capable of accessing Veritas AutoSupport or sending an emailed response to alerts on their own behalf or for other cluster members, if necessary.

Conclusion

With the ever-increasing importance of data and the continuously growing data footprints being managed by organizations today, a resilient, scalable, and efficient data protection strategy is key to providing resilient IT services and reducing operational complexity. NetBackup Flex Scale is a highly flexible solution that provides several key benefits focused on enterprise data protection:

- A hyperconverged platform to help reduce operational complexity and time to value
- A scale-out architecture capable of self-balancing while adding capacity
- Integrated high availability and resiliency for backup and catalog data
- Data and infrastructure immutability in lockdown mode to guard against ransomware or the malicious deletion or alteration of data
- A familiar user experience with the flexibility to use NetBackup management interfaces or a single API endpoint for custom orchestration

Providing efficient, reliable, and scalable data protection for modern, complex IT environments can be a significant challenge. With multiple applications and systems that have different data protection requirements, organizations are often burdened with inflexible solutions that can be difficult to use and scale. NetBackup Flex Scale solves this problem by providing an easily scalable hyperconverged platform that reduces operational complexity and expands on NetBackup's industry-leading enterprise data protection features and capabilities.

Resources and References

- NetBackup Flex Scale Product Page – https://www.veritas.com/protection/netbackup/netbackup-flex-scale
- NetBackup Flex Scale Product Documentation – https://sort.veritas.com/documents?prod=NBUFS
- NetBackup Flex Scale Data Sheets
  https://www.veritas.com/content/dam/www/en_us/documents/data-sheet/DS_netbackup_flex_scale_V1020
  https://www.veritas.com/content/dam/www/en_us/documents/data-sheet/DS_netbackup_flex_scale_dell_V1505.pdf
- Cohasset Associates – Veritas NetBackup Flex Scale: SEC 17a-4(f), FINRA 4511(c), and CFTC 1.31(c)-(d) Compliance Assessment
  https://www.veritas.com/content/dam/www/en_us/documents/analyst-report/AR_netbackup_flex_scale_cohasset_associates_report.pdf
- NetBackup Flex Scale Technical Overview Video
  https://youtu.be/TrBdZSI-d0I
- NetBackup Flex Scale – Secure by Default white paper
  https://www.veritas.com/content/dam/www/en_us/documents/white-papers/WP_netbackup_flex_scale_secure_by_default_V1517.pdf

About Veritas

Veritas Technologies is a global leader in data protection and availability. Over 80,000 customers—including 87 percent of the Fortune Global 500—rely on us to abstract IT complexity and simplify data management. The Veritas Enterprise Data Services Platform automates the protection and orchestrates the recovery of data everywhere it lives, ensures 24/7 availability of business-critical applications, and provides enterprises with the insights they need to comply with evolving data regulations. With a reputation for reliability at scale and a deployment model to fit any need, Veritas Enterprise Data Services Platform supports more than 800 different data sources, over 100 different operating systems, more than 1,400 storage targets, and more than 60 different cloud platforms. Learn more at www.veritas.com. Follow us on Twitter at @veritastechllc.

VERITAS

2625 Augustine Drive
Santa Clara, CA 95054
+1 (866) 837 4827
veritas.com

For global contact information visit:
veritas.com/company/contact

V1464 03/22